arXiv:1608.06608v1 [cs.LG] 23 Aug 2016

Infinite-Label Learning with Semantic Output Codes

Yang Zhang
Center for Research in Computer Vision
University of Central Florida
Orlando, FL 32816
yangzhang@knigths.ucf.edu

Rupam Acharyya
Department of Computer Science
University of Rochester
Rochester, NY 14627
racharyy@cs.rochester.edu

Ji Liu
Department of Computer Science
University of Rochester
Rochester, NY 14627
jliu@cs.rochester.edu

Boqing Gong
Center for Research in Computer Vision
University of Central Florida
Orlando, FL 32816
bgong@crcv.ucf.edu

August 24, 2016

Abstract

We develop a new statistical machine learning paradigm, named *infinite-label learning*, to annotate a data point with more than one relevant labels from a candidate set, which pools both the finite labels observed at training and a potentially infinite number of previously unseen labels. The infinite-label learning fundamentally expands the scope of conventional multi-label learning, and better models the practical requirements in various real-world applications, such as image tagging, ads-query association, and article categorization. However, *how can we learn a labeling function that is capable of assigning to a data point the labels omitted from the training set?* To answer the question, we seek some clues from the recent work on zero-shot learning, where the key is to represent a class/label by a vector of semantic codes, as opposed to treating them as atomic labels. We validate the infinite-label learning by a PAC bound in theory and some empirical studies on both synthetic and real data.

1 Introduction

We formalize a new machine learning paradigm, infinite-label learning, where the goal is to learn a labeling function to annotate a data point with all relevant labels from a candidate set, which pools both the finite labels observed at training and a potentially infinite number of previously unseen labels. Infinite-labeling learning is motivated by a plethora of real-world applications. As below, we give a few examples to facilitate our further discussion.

Image tagging. More than 50% of the Flickr images have no text tags and are never retrieved for text queries. As a result, automatically suggesting relevant tags/labels for the images is a pressing need and yet a daunting challenge; there are more than 53M tags on Flickr. Moreover, tons of new tags emerge over time.

Ads-query association. Search engines allow advertisers to bid on user queries to deliver ads to targeted audience. It is both commercially intriguing and scientifically interesting to automatically associate the virtually infinite number of user queries (labels) to the ads (data points).

Text categorization. The WIKIPEDIA pages are grouped by more than 200K categories. A text might be about any of science, politics, finance, and/or other topics. Infinite-label learning can save the writers from the tedious process of manually classifying the future WIKIPEDIA pages, even if new categories are added to WIKIPEDIA.

The above problems are actually often used to benchmark different (extreme) multi-label learning algorithms [21, 23, 3]. However, we contend that the conventional multi-label learning tackles the problems on the surface, and fails to consider any new labels that appear in the test stage and were omitted from the training stage. In a sharp contrast, *infinite*-label learning fundamentally expands the scope of *multi*-label learning; it is expected to explicitly handle novel labels that show up after the labeling function has been learned.

How can we learn a labeling function that is capable of assigning to a data point the labels omitted from the training set? To answer the question, we seek some clues from the recent work on zero-shot learning [15, 19, 8, 17], where the key is to represent a class by a vector of semantic codes [15], as opposed to treating them as atomic labels. We can thus learn a labeling function to take as input a data point and a label's semantic codes, and output whether or not the encoded label is relevant to the data. Both seen and unseen labels are encoded by the same mechanism and usually by the same knowledge base. As a result, the labeling function learned from a finite training set is able to extrapolate to the labels unseen during the training. To verify this intuition, we provide a PAC bound (Section 3) and some empirical studies (Section 4) on both synthetic and real data.

The semantic codes of the class labels are often derived from a large knowledge base in addition to the training set. Thanks to the studies in linguistics, words in speech recognition are represented by combinations of phonemes [22]. In visual object recognition, a class is described by a set of visual attributes [10, 7]. More recently, the distributed representations of English words [12, 16] have found their applications in a variety of tasks.

We organize the remaining of the paper as follows. Section 2 formally states the problem of infinite-label learning and contrasts it to some related work. A PAC bound and some theoretical analyses are presented in Section 3. In Section 4, we experimentally study the proposed infinite-label learning and reveal some of its empirical properties. Section 5 concludes the paper.

2 Problem statement, modeling, and algorithmic solutions

In this section, we formally state the problem of infinite-label learning, followed by a modeling assumption about its corresponding data generation process. We also provide some exemplar algorithmic solutions. Finally, we contrast infinite-label learning to the closely related multi-label learning and zero-shot learning.

2.1 Problem statement

Suppose we have a set of semantic codes of the labels $\mathcal{L} = \{\boldsymbol{\lambda}_l \in \mathbb{R}^n\}_{l=1}^{\mathsf{L}}$ and a training sample $\mathcal{S} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m) \in \mathbb{R}^d \times \{-1, +1\}^{\mathsf{L}}\}_{m=1}^{\mathsf{M}}$, where the annotation $y_{ml} = +1$ indicates that the l-th label, which is semantically encoded as $\boldsymbol{\lambda}_l$, is relevant to the m-th data point $\boldsymbol{x}_m$. The semantic codes could be the distributed representations of English words, phoneme composition of words in speech, co-occurrence vector, or visual attributes of objects. For brevity, we refer to the semantic codes of the labels $\mathcal{L} = \{\boldsymbol{\lambda}_l\}_{l=1}^{\mathsf{L}}$ as labels in the rest of the paper. Denote by $\mathcal{U} = \{\boldsymbol{\lambda}_l\}_{l>\mathsf{L}}$ another set of labels disjoint from $\mathcal{L}$. The labels in $\mathcal{U}$ correspond to no data points in the training sample and thus often termed *zero-shot* labels [10, 24]. The objective of **infinite-label learning** is to infer a labeling function $\boldsymbol{h} : \mathbb{R}^d \mapsto 2^{\mathcal{L} \cup \mathcal{U}}$ from the training sample $\mathcal{S}$ and observed labels $\mathcal{L}$, such that $\boldsymbol{h}(\boldsymbol{x}) \subseteq \mathcal{L} \cup \mathcal{U}$ assigns all the relevant labels to a test instance $\boldsymbol{x}$.

2.2 A modeling assumption

We consider the following distributions over the data in infinite-label learning,

$$\boldsymbol{x} \sim P_X, \quad \boldsymbol{\lambda} \sim P_\Lambda, \quad y(\boldsymbol{\lambda}; \boldsymbol{x}) \sim P_{Y|X\Lambda}, \quad \text{where } P_{XY\Lambda} = P_X P_\Lambda P_{Y|X\Lambda}, \tag{1}$$

where $P_{XY\Lambda}$ denotes the joint distribution of the data point $\boldsymbol{x}$, label assignment indicators $\boldsymbol{y}$, and the label $\boldsymbol{\lambda}$. In this paper, we consider a deterministic/noiseless label assignment function $y(\boldsymbol{\lambda}; \boldsymbol{x}) \sim P_{Y|X\Lambda}$, and leave the more generic case to the future work.

Note that we explicitly treat the labels $\mathcal{L} \cup \mathcal{U}$ as an i.i.d. sample drawn from the marginal distribution P_Λ. Moreover, we make the following observation: any data point $\boldsymbol{x}$ can be regarded as incurring a binary classification rule, denoted by $y(\cdot; \boldsymbol{x})$, over the label space, such that $y(\boldsymbol{\lambda}; \boldsymbol{x}) = +1$ tells that the label $\boldsymbol{\lambda} \sim P_\Lambda$ is relevant to the data $\boldsymbol{x}$ and $y(\boldsymbol{\lambda}; \boldsymbol{x}) = -1$ means the opposite. This rule is only partially observed at training in the form of the data annotations $\boldsymbol{y}_m$.

Immediately following the above modeling assumption, our learning objective is to infer the (conditional) classification rule $y(\cdot; \boldsymbol{x})$ for the labels from the training sample $\mathcal{S}$. We consider a fixed set of hypotheses $\mathcal{H}$ for the rule $y(\cdot; \boldsymbol{x})$, and then select a hypothesis $h \in \mathcal{H}$ such that it gives rise to the smallest generalization risk,

$$R(h) = \mathbb{E}_{\boldsymbol{x} \sim P_X} \mathbb{E}_{\boldsymbol{\lambda} \sim P_\Lambda} \mathbb{E}_{y \sim P_{\boldsymbol{Y}|X\Lambda}} [h(\boldsymbol{\lambda}; \boldsymbol{x}) \neq y] \tag{2}$$

where $[\cdot]$ is the 0-1 loss. We provide a PAC bound in Section 3 for the risk along with its empirical estimate from the training set $\mathcal{S}$,

$$\widehat{R}(h) = \frac{1}{\mathsf{ML}} \sum_{m=1}^{\mathsf{M}} \sum_{l=1}^{\mathsf{L}} L(h(\boldsymbol{\lambda}_l; \boldsymbol{x}_m), y_{ml}) + \Omega(h), \tag{3}$$

Table 1: *Multi*-label learning *vs. Infinite*-label learning

	Multi-label learning	Infinite-label learning
Data distribution	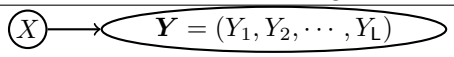	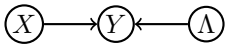
Training data	$\mathcal{S} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m)\}_{m=1}^{\mathsf{M}}$	$\mathcal{S},\ \mathcal{L} = \{\boldsymbol{\lambda}_l\}_{l=1}^{\mathsf{L}}$
Labeling function	$\boldsymbol{f}(\boldsymbol{x}) \subseteq \mathcal{L}$	$\boldsymbol{h}(\boldsymbol{x}) \subseteq \mathcal{L} \cup \mathcal{U},\ \mathcal{U} = \{\boldsymbol{\lambda}_l\}_{l>\mathsf{L}}$
Example 1	$f_l(\boldsymbol{x}) = \operatorname{sgn} \langle \boldsymbol{w}_l, \boldsymbol{x} \rangle,\ 1 \le l \le \mathsf{L}$	$h_l(\boldsymbol{x}) = h(\boldsymbol{\lambda}_l; \boldsymbol{x}) = \operatorname{sgn} \langle V\boldsymbol{x}, \boldsymbol{\lambda}_l \rangle,\ l \ge 1$

where $L(\cdot)$ is a 0-1 loss (in practice one uses a surrogate differentiable loss), and $\Omega(h)$ is a regularization over h. At the high level, the risk of infinite-label learning can be understood as the traditional binary classification risk over the label $\boldsymbol{\lambda}$ and assignment $y|\boldsymbol{x}$, conditioning on the variable X representing the input data.

It is worth pointing out the implication of eq. (3): we employ a "non-canonical" data sampling process from the underlying distribution $P_{XY\Lambda} = P_X P_\Lambda P_{Y|X\Lambda}$. Instead of sampling i.i.d. triplets $\{(\boldsymbol{x}_k, \boldsymbol{\lambda}_k, y_k)\}_{k\ge1}$ from $P_{XY\Lambda}$, we draw from P_X for M times to have $\{\boldsymbol{x}_m\}_{m=1}^{\mathsf{M}}$, draw from P_Λ for L times and obtain $\{\boldsymbol{\lambda}_l\}_{l=1}^{\mathsf{L}}$. After that, we use the (deterministic) label assignment function $y(\boldsymbol{\lambda}; \boldsymbol{x})$ to assign the binary (ir)relevance to all the pairwise combination of data and labels, i.e., we arrive at $\mathsf{M} \times \mathsf{L}$ training data points $\{(\boldsymbol{x}_m, \boldsymbol{\lambda}_l, y_{ml})\}_{m\in[\mathsf{M}], l\in[\mathsf{L}]}$ and they are non-i.i.d.

2.3 Algorithmic solutions

Given the modeling in Section 2.2, the solution to infinite-label learning boils down to how to specify the labeling function $\boldsymbol{h}(\boldsymbol{x}) = \{h(\boldsymbol{\lambda}; \boldsymbol{x}) | \boldsymbol{\lambda} \in \mathcal{L} \cup \mathcal{U}\}$. We give some examples below, and will examine **Example 1** both theoretically (Section 3) and empirically (Section 4).

Example 1. Assuming that each label $\boldsymbol{\lambda} \in \mathcal{L} \cup \mathcal{U}$ is represented by a vector (e.g., word vectors, attributes, or click-through statistics) and $h(\boldsymbol{\lambda}; \boldsymbol{x}) = \operatorname{sgn} \langle V\boldsymbol{x}, \boldsymbol{\lambda} \rangle$, we essentially try to find a hyper-plane $V\boldsymbol{x}$ for classifying either seen or unseen labels $\{\boldsymbol{\lambda}\}$ by minimizing the empirical risk in eq. (3).

Example 2. We can define a new hypothesis set $\mathcal{H}$ by replacing $V\boldsymbol{x}$ in Example 1 with a neural network $\textsc{nn}(\boldsymbol{x})$, such that $h(\boldsymbol{\lambda}; \boldsymbol{x}) = \operatorname{sgn} \langle \textsc{nn}(\boldsymbol{x}), \boldsymbol{\lambda} \rangle$.

Example 3. Consider $h(\boldsymbol{\lambda}; \boldsymbol{x}) = \sum_{l=1}^{\mathsf{L}} \alpha_l(\boldsymbol{x}) k(\boldsymbol{\lambda}, \boldsymbol{\lambda}_l)$, where $\alpha_l(\cdot)$ takes as input a data point $\boldsymbol{x}$ and outputs a scalar and $k(\cdot, \cdot)$ is a kernel function over the labels. The kernel is particularly useful when the semantic codes of the labels are in structured forms (e.g., trees, graphs).

2.4 Related work

The modeling of the data, label, and label assignment in infinite-label learning is in sharp contrast to the prevalent modeling assumptions for the traditional classification problems or the closely related multi-label learning problems, none of which induce the labels $\mathcal{L}$ and the label assignment $\boldsymbol{y}$ by separate distributions. We particularly compare the infinite-label learning to multi-label learning [21, 23, 3] in Table 1. Multi-label learning usually assumes that the training sample $\mathcal{S} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m)\}_{m=1}^{\mathsf{M}}$ is drawn i.i.d. from the joint distribution P_{XY}, totally ignoring the distribution of the labels $\mathcal{L}$. Consequently, the multi-label classifiers cannot deal with any unseen labels in $\mathcal{U}$.

We can understand infinite-label learning as a natural extension of zero-shot learning [15, 19, 8, 17]. They both rely on an external knowledge base for encoding the classes/labels, in addition to the training set, in order to extrapolate the learned classifier or labeling function to handle previously unseen classes/labels. Their difference echos that between multi-label and multi-class classifications, namely, infinite-label learning assigns more than one relevant labels to a data point while under zero-shot learning a data point can belong to one and only one class.

The recent works of label embedding [9, 25, 3, 1] and landmark label sampling [2, 4] are particularly related to ours. However, their label representations are inferred from the training set and are not applicable to unseen labels. Moreover, they lack the basic modeling of the labels as introduced in this paper, that the labels $\boldsymbol{\lambda}$ are drawn i.i.d. from the distribution P_Λ. The most recent works [13, 24] can be regarded as specific instantiations of infinite-label learning, while we formally formulate the problem and provide both theoretical and empirical studies as follows. In addition, the pairwise learning [20] is a very interesting framework which encapsulates matrix completion, collaborative filtering, zero-shot and the infinite-label learning. As a result, our analyses apply to pairwise learning when y_h is binary in [20].

3 Infinite-label learning under the PAC learning framework

In this section, we examine the theoretical properties of infinite-label learning under the PAC learning framework. Given the training sample $\mathcal{S} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m) \in \mathbb{R}^d \times \{-1, +1\}^{\mathsf{L}}\}_{m=1}^{\mathsf{M}}$ and the semantic codes of the seen labels $\mathcal{L} = \{\boldsymbol{\lambda}_l \in \mathbb{R}^n\}_{l=1}^{\mathsf{L}}$, the theorem below sheds lights on the numbers of data points and labels that are necessary to achieve a particular level of error. The supplementary materials give tighter bounds under some sparse regularizations over the hypothesis set $\mathcal{H}$.

Theorem 1. *For any $\delta > 0$ and any $h(\boldsymbol{x}, \boldsymbol{\lambda}) \in \mathcal{H} := \{\operatorname{sgn} \langle V\boldsymbol{x}, \boldsymbol{\lambda}\rangle \mid V \in \mathbb{R}^{n\times d}\}$, the following holds with probability at least $1 - \delta$,*

$$\mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x}, \boldsymbol{\lambda}) \neq y] \leq \frac{1}{\mathsf{ML}} \sum_{m=1}^{\mathsf{M}} \sum_{l=1}^{\mathsf{L}} [h(\boldsymbol{x}_m, \boldsymbol{\lambda}_l) \neq y_{m,l}] + 2 \max\left(\sqrt{\frac{8\log(8/\delta) + 8d\log(2e\mathsf{M}/d)}{\mathsf{M}}}, \sqrt{\frac{8\log(8\mathsf{M}/\delta) + 8n\log(2e\mathsf{L}/n)}{\mathsf{L}}} \right). \quad (4)$$

This generalization error bound is roughly $O(1/\mathsf{M} + 1/\mathsf{L})$ if all log terms are ignored. When M and L converge to infinity, the error vanishes. An immediate implication is the learnability of the infinite-label learning: to obtain a certain accuracy on all labels, one does not have to observe all of them in the training phase. Another interesting implication for practice is that M and L should be chosen in the same order, because $\mathsf{M} \gg \mathsf{L}$ (or $\mathsf{M} \ll \mathsf{L}$) would not bring significant improvements to the generalization bound. Acute readers may notice that the classifier $\langle V\boldsymbol{x}, \boldsymbol{\lambda}\rangle$ can be cast into $\langle V, \boldsymbol{\lambda}\boldsymbol{x}^\top\rangle$, which is in the form of the vanilla linear classifier[1]. However, a fundamental difference from the linear classifier lies on that the sample triplets $(\boldsymbol{x}_m, \boldsymbol{\lambda}_l, y_{ml})$ are not independent in the empirical risk (cf. eq. (3)), causing the main challenge in analysis.

[1] *Remark:* We note that the bound is indeed close to that for the linear classifier except the term of $\log(\mathsf{M})/\mathsf{L}$. This term is likely an artifact of our proof and may be removed by finer analysis.

Proof. To prove the theorem, we essentially need to consider the following probability bound:

$$\begin{aligned}
&\mathbb{P}\left(\exists\ h\in\mathcal{H}:\ \mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{ML}}\sum_{m=1}^{\mathsf{M}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda}_l)\neq y_{m,l})]\geq\epsilon\right)\\
=&\mathbb{P}\Bigg(\exists\ h\in\mathcal{H}:\ \mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]\\
&+\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{ML}}\sum_{m=1}^{\mathsf{M}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda}_l)\neq y_{m,l}]\geq\epsilon\Bigg)\\
\leq&\mathbb{P}\left(\exists\ h\in\mathcal{H}:\ \mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]\geq\frac{\epsilon}{2}\right) \qquad (5)\\
&+\mathbb{P}\left(\exists\ h\in\mathcal{H}:\ \frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{ML}}\sum_{m,l}[h(\boldsymbol{x}_m,\boldsymbol{\lambda}_l)\neq y_{m,l}]\geq\frac{\epsilon}{2}\right) \qquad (6)
\end{aligned}$$

where the last inequality is due to the fact $\mathbb{P}(a+b\geq\epsilon)\leq\mathbb{P}(a\geq\epsilon/2)+\mathbb{P}(b\geq\epsilon/2)$. Next we consider the bound for the two terms (5) and (6) above respectively.

For the first term (5), we have

$$\begin{aligned}
&\mathbb{P}\left(\exists h\in\mathcal{H}:\ \mathbb{E}_{\boldsymbol{x},\boldsymbol{\lambda},y}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{\boldsymbol{\lambda},y|\boldsymbol{x}_m}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]\geq\epsilon_1\right)\\
=&\mathbb{P}\left(\exists h\in\mathcal{H}:\ \mathbb{E}_{\boldsymbol{x},\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\mathbb{E}_{\boldsymbol{\lambda}}\left(\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]\right)\geq\epsilon_1\right)\\
\leq&\mathbb{P}\left(\exists h\in\mathcal{H},\exists\boldsymbol{\lambda}\in\mathcal{L}:\ \mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]-\left(\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]\right)\geq\epsilon_1\right) \qquad (7)
\end{aligned}$$

where the first equality is due to the dependence between $\boldsymbol{\lambda}$ and $\boldsymbol{x}_m$.

Denote $\mathbb{E}_{y|\boldsymbol{x}_m,\boldsymbol{\lambda}}[h(\boldsymbol{x}_m,\boldsymbol{\lambda})\neq y]$ by a function $f(\boldsymbol{x}_m)$. We have $\mathbb{E}(f(\boldsymbol{x}_m))=\mathbb{E}_{\boldsymbol{x}}\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]$. $f(\boldsymbol{x}_1),\cdots,f(\boldsymbol{x}_m)$ are i.i.d. random variables in $\{0,1\}$ — recall that $P_{Y|X\Lambda}$ is deterministic. Define $\mathcal{F}$ to be a hypothesis space

$$\mathcal{F}:=\{f(\boldsymbol{x})=\mathbb{E}_{y|\boldsymbol{x},\boldsymbol{\lambda}}[h(\boldsymbol{x},\boldsymbol{\lambda})\neq y]\mid h\in\mathcal{H},\ \boldsymbol{\lambda}\in\mathcal{L}\}.$$

Using the new notation, we can rewrite (7) in the following

$$\mathbb{P}\left(\exists f\in\mathcal{F}:\ \mathbb{E}(f(\boldsymbol{x}))-\frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}f(\boldsymbol{x}_m)\geq\epsilon_1\right)\leq 4r(\mathcal{F},2\mathsf{M},\mathcal{X})\times\exp(-\mathsf{M}\epsilon_1^2/8) \qquad (8)$$

where the inequality uses the Growth function bound, and $r(\mathcal{F},2\mathsf{M},\mathcal{X})$ is number of maximally possible configurations (or values of $\{f(\boldsymbol{x}_1),\cdots,f(\boldsymbol{x}_{2\mathsf{M}})\}$) given $2\mathsf{M}$ samples in $\mathcal{X}$ over the hypothesis space $\mathcal{F}$. Since $h(\boldsymbol{x},\boldsymbol{\lambda})$ is in the form of $\mathrm{sgn}\langle V\boldsymbol{x},\boldsymbol{\lambda}\rangle$, $r(\mathcal{F},2\mathsf{M},\mathcal{X})$ is bounded by $r(\mathcal{W},2\mathsf{M},\mathcal{X})$ with $\mathcal{W}:=\{\mathrm{sgn}(\boldsymbol{w}^\top\boldsymbol{x})\mid\boldsymbol{w}\in\mathbb{R}^d\}$. Hence, we have

$$r(\mathcal{F},2\mathsf{M},\mathcal{X})\leq r(\mathcal{W},2\mathsf{M},\mathcal{X})\leq\left(\frac{2e\mathsf{M}}{d}\right)^{\mathbf{VC}(\mathbb{R}^d)}=\left(\frac{2e\mathsf{M}}{d}\right)^d \qquad (9)$$

where $\mathbf{VC}(\mathbb{R}^d)$ is the VC dimension for the space $\mathbb{R}^d$. We then have

$$(8) \leq 4\left(\frac{2e\mathsf{M}}{d}\right)^d \times \exp(-\mathsf{M}\epsilon_1^2/8).$$

By letting the right hand side equal $\delta/2$, we have

$$\epsilon_1 = \sqrt{\frac{8\log 8/\delta + 8d\log(2e\mathsf{M}/d)}{\mathsf{M}}}. \tag{10}$$

Next we consider the upper bound for (6):

$$\begin{aligned}
&\mathbb{P}\left(\exists h \in \mathcal{H} \ : \ \frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}} \mathbb{E}_{\boldsymbol{\lambda}, y|\boldsymbol{x}_m}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}) \neq y] - \frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\frac{1}{\mathsf{L}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}_l) \neq y_{m,l}] \geq \epsilon_2\right) &(11)\\
=&\mathbb{P}\left(\exists h \in \mathcal{H} \ : \ \frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}} \mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x}_m, \boldsymbol{\lambda}}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}) \neq y] - \frac{1}{\mathsf{M}}\sum_{m=1}^{\mathsf{M}}\frac{1}{\mathsf{L}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}_l) \neq y_{m,l}] \geq \epsilon_2\right)\\
\leq&\mathbb{P}\left(\exists h \in \mathcal{H},\ \exists m \in \{1, \cdots, \mathsf{M}\} : \ \mathbb{E}_{\boldsymbol{\lambda}}\mathbb{E}_{y|\boldsymbol{x}_m, \boldsymbol{\lambda}}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}) \neq y] - \frac{1}{\mathsf{L}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_m, \boldsymbol{\lambda}_l) \neq y_{m,l}] \geq \epsilon_2\right)\\
\leq&\mathsf{M} \times \mathbb{P}\left(\exists h \in \mathcal{H} : \ \mathbb{E}_{y, \boldsymbol{\lambda}|\boldsymbol{x}_1}[h(\boldsymbol{x}_1, \boldsymbol{\lambda}) \neq y] - \frac{1}{\mathsf{L}}\sum_{l=1}^{\mathsf{L}}[h(\boldsymbol{x}_1, \boldsymbol{\lambda}_l) \neq y_{1,l}] \geq \epsilon_2\right). &(12)
\end{aligned}$$

where the last inequality uses the union bound. Denote $[h(\boldsymbol{x}_1, \boldsymbol{\lambda}) \neq y]$ by $g(\boldsymbol{\lambda}, y)$ for short. Then we have $\mathbb{E}(g(\boldsymbol{\lambda}, y)) = \mathbb{E}_{y, \boldsymbol{\lambda}|\boldsymbol{x}_1}[h(\boldsymbol{x}_1, \boldsymbol{\lambda}) \neq y]$ and $g(\boldsymbol{\lambda}_1, y_{1,1}), \cdots, g(\boldsymbol{\lambda}_L, y_{1,L})$ are i.i.d. samples taking values from $\{0, 1\}$ given $\boldsymbol{x}_1$. Define $\mathcal{G}$ to be a hypothesis space

$$\mathcal{G} := \{g(\boldsymbol{\lambda}, y) = [h(\boldsymbol{\lambda}, \boldsymbol{x}_1) \neq y] \mid h \in \mathcal{H}\}.$$

Then we can cast the probabilistic factor in (12) into

$$\mathbb{P}\left(\exists g \in \mathcal{G} : \ \mathbb{E}(g(\boldsymbol{\lambda}, y)) - \frac{1}{\mathsf{L}}\sum_{l=1}^{\mathsf{L}} g(\boldsymbol{\lambda}_l, y_l) \geq \epsilon_2\right) \leq 4r(\mathcal{G}, 2\mathsf{L}, \mathcal{L}) \times \exp(-\mathsf{L}\epsilon_2^2/8) \tag{13}$$

where the last inequality uses the Growth function bound again. To bound $r(\mathcal{G}, 2\mathsf{L}, \mathcal{L})$, we consider the structure of $h(\boldsymbol{x}, \boldsymbol{\lambda}) = \mathrm{sgn}(V\boldsymbol{x}, \boldsymbol{\lambda})$ and define $\mathcal{W}' = \{\mathrm{sgn}(\boldsymbol{z}^\top\boldsymbol{\lambda}) \mid \boldsymbol{z} \in \mathbb{R}^n\}$. It suggests that

$$r(\mathcal{G}, 2\mathsf{L}, \mathcal{L}) \leq r(\mathcal{W}', 2\mathsf{L}, \mathbb{R}^n) \leq \left(\frac{2e\mathsf{L}}{n}\right)^{\mathbf{VC}(\mathbb{R}^n)} = \left(\frac{2e\mathsf{L}}{n}\right)^n.$$

Hence we have

$$(11) \leq 4\mathsf{M}\left(\frac{2e\mathsf{L}}{n}\right)^n \times \exp(-\mathsf{L}\epsilon_2^2/8).$$

By letting the right hand side equal $\delta/2$, we have

$$\epsilon_2 = \sqrt{\frac{8\log 8\mathsf{M}/\delta + 8n\log(2e\mathsf{L}/n)}{\mathsf{L}}}. \tag{14}$$

Plugging (10) and (14) into (5) and (6) respectively, we prove the theorem.

□

4 Empirical studies

In this section, we continue to investigate the properties of infinite-label learning. While the theory result in Section 3 justifies its "learnability" in some sort, there are many other questions of interest to explore for the practical use of infinite-label learning. We focus on the following two questions and provide some empirical insights using respectively synthetic data and real data.

1. After we learn a labeling function from the training set, how many and what types of unseen labels can we confidently handle using this labeling function?
2. What is the effect of varying the number of seen labels $\mathcal{L}$, given a fixed union $\mathcal{L} \cup \mathcal{U}$ of seen and unseen labels? Namely, the same set of labels $\mathcal{L} \cup \mathcal{U}$ will be used at the test stage, but we learn different labeling functions by varying the seen labels $\mathcal{L}$.

4.1 Synthetic experiments

We generate some synthetic data to answer the first question. This offers us the flexibility to control the number of labels.

Data. We randomly sample 500 training data points $\{\boldsymbol{x}_m \in \mathbb{R}^3\}_{m=1}^{\mathsf{M}=500}$ and 1000 testing data points from a five-component Gaussian mixture model. We also sample 10 seen labels $\mathcal{L} = \{\boldsymbol{\lambda}_l \in \mathbb{R}^2\}_{l=1}^{\mathsf{L}=10}$ and additional 2990 unseen labels $\mathcal{U} = \{\boldsymbol{\lambda}_{11}, \cdots, \boldsymbol{\lambda}_{3000}\}$ from a Gaussian distribution. Note that only the seen labels $\mathcal{L}$ are revealed during the training stage. As below specifies the distributions,

$$\boldsymbol{x} \sim P(\boldsymbol{x}) = \sum_{k=1}^{5} \pi_k \mathcal{N}\left(\boldsymbol{\mu}_k, \boldsymbol{U}_k \boldsymbol{U}_k^T, \right), \quad \boldsymbol{\lambda} \sim \mathcal{N}\left(\begin{bmatrix} 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 1 & 1.5 \\ 1.5 & 3 \end{bmatrix}\right), \quad \boldsymbol{x} \in \mathbb{R}^3, \boldsymbol{\lambda} \in \mathbb{R}^2$$

where the mixture weights π_k are sampled from a Dirichlet distribution $\textsc{Dir}(3,3,3,3,3)$, and both mean μ_k and $\boldsymbol{U}_k$ are sampled from a standard normal distribution (using the `randn` function in MATLAB). Finally, we generate a "groundtruth" matrix $V \in \mathbb{R}^{2\times 3}$ from a standard normal distribution. The label assignments are thus given by $y_{ml} = \operatorname{sgn}\langle V\boldsymbol{x}_m, \boldsymbol{\lambda}_l\rangle$ for both training and testing data and both seen and unseen labels. For the training set $\mathcal{S} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m)\}_{m=1}^{\mathsf{M}=500}$, we randomly flip the sign of each y_{ml} with probability $p = 0.1$.

We deliberately choose the low dimensions for the data and labels so that we can visualize them and have a concrete understanding about the results to be produced. Figure 1(a) and (b) show the data points and labels we have sampled. The training data and the seen labels are in red color, while all the other (test) data points and labels are unseen during training.

Given the training set $\mathcal{S}$, we learn the model parameters by minimizing a hinge loss

$$\widehat{V} \leftarrow \arg\min_{V} \frac{1}{\mathsf{ML}} \sum_{m=1}^{\mathsf{M}=500} \sum_{l=1}^{\mathsf{L}=10} \max\left(1 - y_{ml}\left\langle V\boldsymbol{x}_m, \boldsymbol{\lambda}_l\right\rangle, 0\right),$$

and then try to assign both seen and unseen labels $\mathcal{L} \cup \mathcal{S}$ to the 1000 test data points, using $\widehat{V}$. We challenge the learned infinite-label assignment model $y_{ml} = \operatorname{sgn}\left\langle \widehat{V}\boldsymbol{x}_m, \boldsymbol{\lambda}_l \right\rangle$ by gradually increasing the difficulties. Namely, we rank all the labels according to their distances to the seen labels $\mathcal{L}$, where the distance between $\boldsymbol{\lambda}_l$ and the seen labels $\mathcal{L}$ is calculated by $\min_{\boldsymbol{\lambda}\in\mathcal{L}} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_l\|_2$. We then

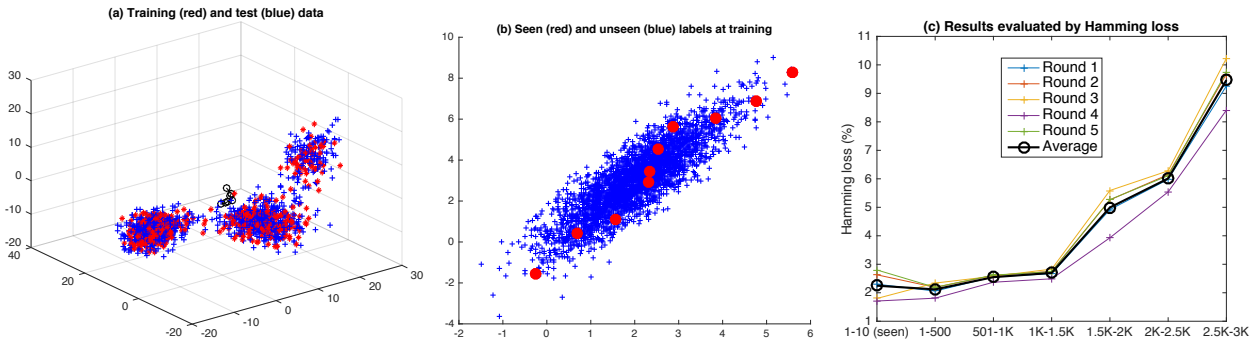

Figure 1: Synthetic data points (a), labels (b), and experimental results given different unseen labels (c). The data and labels used in training are in the red color. We evaluate the infinite-label classification results on the testing data and the unseen labels by the Hamming loss.

evaluate the label assignment results given every 500 consecutive labels in the ranking list (as well as the 10 seen labels). Arguably, the last 500 labels, which are the furthest subset of labels from the seen labels $\mathcal{L}$, impose the biggest challenge to the learned model.

Results. Figure 1(c) shows the label assignment errors for different subsets of labels. We run 5 rounds of experiments each with different randomly sampled data, and report their results as well as the average. We borrow from multi-label classification [23] the Hamming loss as the evaluation metric. It is computed by $\frac{1}{1000}\sum_{m=1}^{1000}(|\boldsymbol{y}_m \neq \boldsymbol{z}_m|^T\mathbf{1})/|\boldsymbol{y}_m|^T\mathbf{1}$, where $\boldsymbol{y}_m$ is the groundtruth label assignments to the data point $\boldsymbol{x}_m$ and $\boldsymbol{z}_m$ is the predictions. Note that this is inherently different from the classification accuracy/error used for evaluating multi-class classification.

We draw the following observations from Figure 1(c). First, infinite-label learning is feasible since we have obtained decent classification results for 3000 labels although only 10 of them are seen in training. Second, when the unseen labels are not far from the seen labels, the label assignment results are on par with the performance of assigning only the seen labels to the test data (cf. the Hamming losses over the first, second, and third 500-label subsets). Third, labels that are far from the seen labels may cause larger confusions to the infinite-label assignment model learned from finite seen labels. Increasing the number of seen labels and/or data points during training can improve the model's generalization capability to unseen labels.

4.2 Image tagging

We experiment with image tagging to seek empirical answers to the second question raised at the beginning of this section.

Image tagging is a real scenario that can benefit from the proposed work. It is often defined as ranking a set of labels for a query image. The main challenge in practice, however, is that the candidate set could be very large—there are about 53M tags on Flickr. Under infinite-label learning, one does not need to see all the candidate tags during training; otherwise much tedious human efforts are required to tag the images. Instead, we can learn a labeling function from a small number of seen tags and then apply it to all the possible tags. In this case, how many seen tags should we use in order to achieve about the same results as using all the tags for training? This is exactly the second question we asked at the beginning of this section.

Data. We conduct our experiments using the NUS-WIDE dataset [6]. It has 269,648 images in total although we are only able to retrieve 223,821 images using the provided image URLs. Among them, 134,281 are training images and the rest are testing images, according to the split by NUS-WIDE. We further randomly choose 20% of the training images as a validation set. The image features $\boldsymbol{x}$ are l_2-normalized VGGNet-19 [18] last fully-connected-layer activations.

Each image in NUS-WIDE has been manually annotated with the relevant tags out of 81 candidate tags in total. We obtain their semantic codes by the pre-trained GloVe word vectors [16]. While all the 81 tags are considered at the test stage, we randomly choose $\mathsf{L} = 81$ (100% out of the 81 tags), 73 (90% out of the 81 tags), 65 (80% out of 81), 57 (70% out of 81), 49 (60% out of 81), and 41 (50% out of 81) seen tags for training different labeling functions $h(\boldsymbol{\lambda}; \boldsymbol{x}) = \operatorname{sgn} \langle V\boldsymbol{x}, \boldsymbol{\lambda} \rangle$. Note that some training images would have no relevant tags under some settings; we simply drop them out from the training and validation sets.

Learning and evaluation. Image tagging is often evaluated based on the top few tags returned by a system, assuming that users do not care about the remaining tags. We report the results measured by four popular metrics: Mean image Average Precision (MiAP) [11] and the top-3 precision, recall, and F1-score. Accordingly, in order to impose the ranking property to our labeling function, we learn it using the RankNet loss [5],

$$\widehat{V} \leftarrow \arg\min_{V} \frac{1}{\mathsf{M}} \sum_{m=1}^{\mathsf{M}} \frac{1}{(\mathsf{L} - \mathsf{K}_m)\mathsf{K}_m} \sum_{\substack{k:y_{mk}=+1 \\ \bar{k}:y_{m\bar{k}}=-1}} \log(1 + \exp(\langle V\boldsymbol{x}_m, \boldsymbol{\lambda}_{\bar{k}} \rangle - \langle V\boldsymbol{x}_m, \boldsymbol{\lambda}_k \rangle)) + \gamma \left\| V \right\|_2 ,$$

where K_m is the number of relevant tags to the m-th image. The hyper-parameter γ is tuned by the MiAP results on the validation set for each experiment.

Baselines. We compare our results to those of two zero-shot learning methods, respectively proposed by Norouzi et al. [14] and Akata et al. [1]. Both methods are developed for multi-class classification, namely, to classify an image to one and only one of the tags/classes. When we use them to solve the multi-label image tagging problem, different tags/classes will have duplicated images. In other words, they are naturally flawed in solving the infinite-label learning problem. We still include them as the references here, since there is no prior work to handle both multi-label learning and unseen labels at testing, to the best of our knowledge.

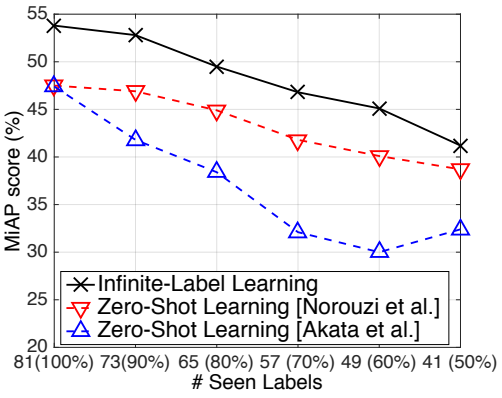

Figure 2: The infinite-label learning results for image tagging under different numbers of seen labels.

Results. Figure 2 shows the MiAP results under different numbers of seen tags. We include the results evaluated by other metrics (the top-3 precision, recall, and F1-score) in the supplementary materials. Recall that although we train the labeling functions using different numbers of seen

tags, the task at the testing stage remains the same for them: the labeling functions rank all the 81 tags for each test image (and return the top few as the suggested tags).

We can see that, without any surprise, the performance decreases as less seen tags are used for training. However, the decrease is fairly gentle—the MiAP drops by 1% and 3% (absolutely) using from 100% to 90% and from 90% to 80% of the 81 tags for training, respectively. Besides, no matter under which experiment setting, the proposed infinite-label learning outperforms the zero-shot learning methods by a large margin, thanks to that it is a more appropriate solution to the image tagging problem. It is interesting to see that the gaps between infinite-label learning and zero-shot learning methods are smaller as the seen tags become less, likely because the overlapped images between different tags/classes are reduced and the tagging problem approaches multi-class classification in training.

5 Conclusion

In this paper, we have proposed a new machine learning paradigm which we term infinite-label learning. It fundamentally expands the scope of multi-label learning, in that at the testing stage, the learned labeling function can assign potentially an infinite number of relevant labels to a data point. Infinite-label learning is made feasible by representing the labels using semantic codes. We provide a full treatment to the infinite-label learning, including a new modeling assumption about the data generation process, a PAC bound, and empirical studies about its performance and properties.

There are many avenues for the future work to further explore infinite-label learning. We note that our current bound can be likely improved. Theoretical understanding about its performance under the MiAP evaluation is also necessary given that MiAP is prevalent in evaluating the multi-label results. One particularly interesting application of infinite-label learning is on the extreme multi-label classification problems [3].

References Cited

[1] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2927–2936, 2015.

[2] Krishnakumar Balasubramanian and Guy Lebanon. The landmark selection method for multiple output prediction. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 983–990, 2012.

[3] Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain. Sparse local embeddings for extreme multi-label classification. In *Advances in Neural Information Processing Systems*, pages 730–738, 2015.

[4] Wei Bi and James Kwok. Efficient multi-label classification with many labels. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 405–413, 2013.

[5] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96. ACM, 2005.

[6] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, page 48. ACM, 2009.

[7] Alireza Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1778–1785. IEEE, 2009.

[8] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems*, pages 2121–2129, 2013.

[9] Daniel Hsu, Sham Kakade, John Langford, and Tong Zhang. Multi-label prediction via compressed sensing. In *NIPS*, volume 22, pages 772–780, 2009.

[10] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(3):453–465, 2014.

[11] Xirong Li, Tiberio Uricchio, Lamberto Ballan, Marco Bertini, Cees GM Snoek, and Alberto Del Bimbo. Socializing the semantic gap: A comparative survey on image tag assignment, refinement and retrieval. *arXiv preprint arXiv:1503.08248*, 2015.

[12] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

[13] Jinseok Nam, Eneldo Loza Mencía, Hyunwoo J Kim, and Johannes Fürnkranz. Predicting unseen labels using label hierarchies in large-scale multi-label learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 102–118. Springer, 2015.

[14] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650*, 2013.

[15] Mark Palatucci, Dean Pomerleau, Geoffrey E Hinton, and Tom M Mitchell. Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, pages 1410–1418, 2009.

[16] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543, 2014.

[17] Bernardino Romera-Paredes and PHS Torr. An embarrassingly simple approach to zero-shot learning. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 2152–2161, 2015.

[18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[19] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *Advances in neural information processing systems*, pages 935–943, 2013.

[20] Michiel Stock, Tapio Pahikkala, Antti Airola, Bernard De Baets, and Willem Waegeman. Efficient pairwise learning using kernel ridge regression: an exact two-step method. *arXiv preprint arXiv:1606.04275*, 2016.

[21] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Mining multi-label data. In *Data mining and knowledge discovery handbook*, pages 667–685. Springer, 2009.

[22] Alex Waibel. Modular construction of time-delay neural networks for speech recognition. *Neural computation*, 1(1):39–46, 1989.

[23] Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *Knowledge and Data Engineering, IEEE Transactions on*, 26(8):1819–1837, 2014.

[24] Yang Zhang, Boqing Gong, and Mubarak Shah. Fast zero-shot image tagging. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

[25] Yi Zhang and Jeff G Schneider. Multi-label output codes using canonical correlation analysis. In *International Conference on Artificial Intelligence and Statistics*, pages 873–882, 2011.